2. The Opportunities

The international community has a long and effective record of framing and reaching agreement in almost all areas of interaction. These are especially powerful in areas of standards, quality controls, certifications and so forth. As a result, we should take stock of what we do know about what works best in different areas and domains.

Furthermore, how and why do we know what works best? These questions are designed to empower researchers, businesses, government agencies, and international entities – private and public – to address how and why? Then, too, given the known "unknowns," what should we know? We have an opportunity to mine our own records for the "best fit" with the properties and conditions surrounding the current Artificial Intelligence dilemma.

Among the major opportunities before us is to inquire: What is the best precedent? Is it nuclear power? Is it climate change? What are other high-risk areas? Usually, we respond to such questions long after the fact. But can we avoid this delay? At this point, we have an opportunity to seriously consider the properties of a global accord in AI before we are faced with a major disaster.

Of high value, for example, is to consider and address the role of ethics in courses on innovations in AI, as well as ethics for all uses and users. So, too, it is important to focus on international law relevant to AI. There are many other high-value issues to consider at this point. The reason is this: The lines of political contention are not yet clearly drawn among potentially conflicting perspectives (or countries).

Therefore, now is the opportunity to proceed before these are consolidated into "lines in the sand."

3. The Next Steps

At this point governments do not control AI innovation and/or diffusion. Much of the action taken is not in the public sector. Individuals and non-state groups matter and matter a lot. Constituencies are varied and overlapping. Consensus building is essential for society, not only for governments. Any position taken must be in the interest of everyone. Any initiative cannot be seen to dampen innovation or markets.

At the same time, we know from experience that "punishments"—in their various forms, do not work. We are in a world where large firms in the IT and communication business area are very powerful. Many are larger than most countries. The dilemma becomes: Whom to punish?

The immediate next step is establishing a multi-stakeholder support base. This is a necessary step to get to the point of articulation of interests and negotiations for "best outcomes."

We are now dealing with 21st century realities wherein state coalition building is essential. We could even initiate a global competition among young minds for creating the best international agreement on artificial intelligence.

03

An AI Ecosystem of Trust

URSULA VON DER LEYEN
PRESIDENT OF EUROPEAN COMMISSION

03

An AI Ecosystem of Trust

URSULA VON DER LEYEN
PRESIDENT OF EUROPEAN
COMMISSION

(From a speech to the Boston Global Forum, December 12, 2020)

Dear Governor Dukakis, Ladies and Gentlemen, It is such an honour to be here with you today. At the Michael Dukakis Institute for Leadership and Innovation, you are at the forefront of research and debate. And you definitely work on some of the world's most pressing issues. You drive the discussion on digital policy and how a human-centric approach on AI could look like. This is an issue whose importance simply cannot be overestimated. Today I would like to speak about our European perspective.

I am a tech optimist! I believe that science and technology have the power to improve our lives. When I studied to become a medical doctor, I learned that technology saves lives. Artificial intelligence can help identify cancers that used to go undetected. And modern robots can help to perform high-precision surgery that used to be too risky just a few years ago.

But of course, there is another side. Some of us are deeply concerned about the role which will be left for human beings in a world run by AI. Others worry about the serious effects that algorithms can have on the health of our democracies. Who is taking the final decisions? Who is steering the flow of information? Who is deciding on rules?

The world that we see through social media platforms seems real. But it is constructed. A list of search results seems objective. But it is different for each and every one of us. Steered by algorithms. Thousands of likes and retweets create a sense of unity. While we literally live in different worlds.

Yes, algorithms can be a danger to our democracy. But they do not have to be. We have the power to protect ourselves. This is why, next week, the European Commission will present its Digital Services Act to make sure that platforms are transparent about the way algorithms work. And that they take responsibility for their systemic effects. We just cannot leave decisions, which have a huge impact on our democracies, to systems, which are a black box for us. There must be at least transparency on how the algorithm works.

AI can have profound impacts on the life of the individual. AI may influence who to recruit for a certain post or whether to grant a certain pension application. For people to accept a role for AI in such decisions, they must be comprehensible. And they must respect people's legal rights – just like any human decision-maker must. This is why we have to be able to examine the workings of the system and to ensure human oversight. Our aim is to create an AI ecosystem of trust.

This is our European way of dealing with AI: Yes, Europe embraces innovation and entrepreneurship. And yes, we are eager to explore the full potential of AI for our industry and services. But we will never lose sight of those who are meant to benefit from these technological wonders – our citizens. What sets Europe apart from competitors like China is not the size of our tech sector or how attractive we are for venture capital. What sets Europe apart from competitors is the fact that our values come first. Human beings come first.

I am so honoured to receive this award today. For me this is also a very personal starting point for a great new beginning between the U.S. and Europe. A partnership that has been tested time and again. A partnership that today is needed more than ever, not at least in the digital world. The way we approach algorithms and AI will define the world we live in.

This is why the EU proposes to start work on a Transatlantic AI Agreement. We want to set a blueprint for regional and global standards aligned with our values: Human rights, and pluralism, inclusion and the protection of privacy. A transatlantic dialogue on the responsibility of online platforms!

Together, we could set the blueprint for other democracies facing the same challenges. To make sure that algorithms and AI are a force for good. Once again: I am delighted to be with you today!

To view President of the EU Commission Ursula von der Leyen's speech: https://youtu.be/Cy9jbwT0CEI or https://ec.europa.eu/commission/presscorner/detail/en/speech_20_2402

An AI "Bill of Rights" in the Digital Age

04

STAVROS LAMBRINIDIS
EU AMBASSADOR TO UNITED STATES

(From speech at Boston Global Forum Conference, April 28, 2021)

Dear Governor Dukakis, Dear Tuan Nguyen and Mark Rotenberg, Distinguished Members of the AIWS Committee, Ladies and Gentlemen,

Thank you for this distinction, which I accept with gratitude and humility. I am especially moved to receive it in the presence of Governor Dukakis, who not only epitomizes public service and leadership in this country but has also been a source of great pride in his ancestral one. Σας ευχαριστώ για τη μεγάλη τιμή που μου κάνετε, thank you for this high honor!

Introduction

Ambassadors, as is well known, are asked to speak about many topics for which they are not experts themselves. But in fact, Artificial Intelligence and human rights in the digital age are two topics that I have followed closely for a long time. They are a large portion of my current work as European Union Ambassador to the United States, and long before that, I followed these issues closely as a Greek Member of the European Parliament. While Vice President of the Parliament's civil liberties committee, I wrote the very first report on security and privacy in the digital age, exactly 11 years ago, when the topic was hardly as popular as it is today.

Perhaps my personal experience of growing up under a dictatorship in Greece endowed me with an acute awareness of the fragility of our open and free societies. As a child, I saw first-hand that even strong democracies can fall under authoritarian spells. I remember the dictatorship holding "files" on different citizens (including my parents), with personal information revealing their political activities and preferences, to be used against them, or to scare them into submission or complacency.

Perhaps this explains best why, when I look at the promise but also the challenges posed by digital technologies, I have always been guided by two fundamental principles, in politics and now diplomacy:

First – in real democracies, it is the people who should have the power to judge the thoughts and actions of their governments, and to hold governments and companies to account; not governments or companies who are supposed to observe and judge the daily actions or thoughts of their citizens. If the unrestrained use of technology leads to the latter instead of the former, we will have flipped Democracy on its head.

Second — In today's democracies, a "Big Brother" will materialize slowly and by stealth, not suddenly in the form of an authoritarian figure who takes away our rights in one fell swoop. If it happens, it will be gradual, by a thousand cuts, with our own explicit or tacit "consent," with our complacency.

In the mid 2000's, an example that illustrated the conundrum was the unfolding mass use of cameras in the streets. What should be their proper use? For regulating traffic? Sounds reasonable. Protecting us from terrorist attacks? Sounds reasonable too. But given that by their very nature they could be used on a 24-hour basis for many more things – identifying all participants at a protest in case just a few of them turned violent? Catching a pickpocket, in addition to a terrorist? – the question quickly became, "Where do democracies draw the line for the use of technology to avoid dangerous slippery slopes? What is necessary, appropriate and proportionate usage? Who should have access to the data and who should not? Where should such personal data be stored in order to be kept safe, and when should it be permanently deleted?" And, soon thereafter, similar questions started to be raised on the collection and use of citizens' personal data by major private companies and digital platforms as well.

The argument used by some governments and businesses at the time, the one that most troubled me, was: "If you have nothing to hide, you have nothing to fear." It troubled me, because it in essence encouraged the "innocent" to offer their consent to their very own unfettered surveillance, in the name of catching the few "guilty." If successful, it could indeed lead to the gradual and irreversible salami-slicing of our rights, "with our consent."

So what I would answer to that posed pseudo-dilemma was, "If you have nothing to hide, you don't have a life!" Because in fact, we all have thousands of elements of our private transactions, relationships, histories, or beliefs, all perfectly legal, that we do not wish others to have unrestricted access to.

That was then. The reality today is that we are all thinking about technology, innovation, and privacy quite differently than we did a few years ago.

We've seen — and are seeing more every day— that Americans want baseline privacy protections, that the status quo isn't good enough, and that the time is ripe for new actions to improve citizens' rights and trust in technology.

In Europe, we have always been forward leaning when it comes to the protection of privacy – perhaps for historic reasons – and withstood significant criticism for pioneering it in the beginning, not least of all with the General Data Protection Regulation (GDPR).

Without a doubt, we are a better and stronger Union today for the privacy safeguards we have in place. And we are now committed to seeing that our privacy protection also safeguards innovation and competition. Our economies and societies need both.

Artificial Intelligence

Whether you call it AI or machine learning, both – in the broadest sense – represent change. Change makes many of us uncomfortable, because it creates a new reality, something different from what we have come to know. It is easy to fear the unknown.

I will not lecture you on the textile revolution, the industrial revolution, or the introduction of the automobile. They represented massive shifts in our technological progress, and the economic benefits as well as the social upheaval that accompanied them are well-documented.

Artificial Intelligence is different, more complex, and far-reaching. In creating a tool that can make judgements – that can decide for us between multiple alternatives – we have introduced a new form of change into our daily lives. It is, if you like, change "to the nth power," scaling exponentially in a way we have not yet experienced.

As policymakers, citizens, and consumers – even as ordinary human beings – we must ask ourselves: Who do we want to make the rules for tools that are becoming increasingly embedded, invisibly, in the fabric of our society? How do we ensure that the AI embedded in the cars we drive, the buildings we use, the energy we consume, the health services we receive, the messages we send, the news we read, even in the refrigerators we use – is safe, controllable, unbiased, and trustworthy? That AI does not discriminate, is not used to "observe and judge," or to impinge on our universal human rights?

In the final analysis, how do we ensure that AI technologies enhance and protect our freedoms, our well-being, and our democracies rather than diminish them?

Europe's balance between Innovation and Fundamental Rights

In Europe, we believe that there is a clear interrelation between Innovation and Fundamental Rights – that one can promote the other.

We value, champion, and thrive from innovation. Last year, as the deadly COVID-19 pandemic spread rapidly across the globe, AI demonstrated its potential to aid humanity by helping to predict the geographical spread of the disease, diagnose the infection through computed tomography scans, and develop the first vaccines and drugs against the virus.

European companies and innovators have been at the forefront in every aspect of that effort. The winner of last year's Future Unicorn Award, presented annually by the European Union to start-ups with the greatest potential, was awarded to a Danish company, Corti, which uses AI and voice recognition to help doctors predict heart attacks.

Clearly, the possibilities and opportunities for AI are immense – from turning on wind turbines to produce the clean energy for our green transition, to detecting cyber-attacks faster than any human being, or cancer in mammograms earlier and more reliably than trained doctors. We hope that AI will even help us to detect the next infectious outbreak, before it becomes a deadly pandemic.

We want AI to do all of these great things. At the same time, just as in every technological evolution that has come before it, we must prepare for the unexpected. With the increasing adoption of AI, our rights to privacy, dignity, freedom, equality, and justice are all at stake. These are fundamental to our lives as Europeans and enshrined in the European Charter of Fundamental Rights.

If it is our aspiration to create machines that are able to do more and more of our thinking, selections and decision-making, we must also take care to ensure they do not make the same mistakes that we humans have been prone to make. Let me offer two examples to illustrate the point:

First

The use of facial, voice, and movement recognition systems in public places can help make our lives more secure. However, it can also allow governments to engage in mass surveillance, intimidation, and repression, as China has shown, in the most cynical and calculated way, in Xinjiang.

Second

The use of AI in recruitment decisions can be helpful. However, if a computer compares resumes of senior managers and concludes that being male is a good predictor of success, the data simply reflects bias – a bias within our society, which historically has favored men for leadership positions. We do not want AI to reinforce existing biases by copying and infinitely replicating them.

These are just two examples that illustrate why we must not become bystanders to the development and deployment of AI. If we, the major world democracies, do not move to establish a regulatory framework, if we do not move fast, smart and strategically to build alliances and set standards for human-centric, trustworthy, and human rights-respecting AI with countries big and small from all over the world, I dread to think who might.

EU Proposal for AI

In Europe, we have been thinking about these questions for many years now. We see that technology is an inescapable, necessary and desirable part of our future. But without trust in it, our progress as a society will simply not be sustainable.

Wearing both my Ambassador and European citizen hats, I am immensely proud that the European Commission has just presented a ground-breaking proposal for a regulatory framework on AI. It is the first proposal of its kind in the world, and it builds on years of work, analysis and consultation with citizens, academics, social partners, NGOs, businesses (including U.S. businesses), and EU Member states.

It is not a regulation for regulation's sake. It responds to calls for a comprehensive approach across the European Union to protect basic rights, encourage innovation consistent with our values, provide legal certainty to innovative companies, spur technological leadership, and prevent the fragmentation of our single market.
In terms of the scope, the draft regulation is actually quite limited. It will introduce a simple classification system with four levels of risk – unacceptable, high, limited, and minimal.

"Unacceptable" and thus prohibited AI practices are those which deploy subliminal techniques beyond a human's consciousness, such as toys or equipment using voice assistance that could lead to dangerous behavior, or the exploitation of the vulnerabilities of specific groups of persons due to their age, physical, or mental disability. Real-time remote biometric identification systems used in public spaces are also classified as an "unacceptable risk" – with extremely narrow exceptions when strictly necessary. As are "social scoring" practices, where governments "score" their citizens as opposed to the other way around.

When enacted, the regulation will also set binding requirements for a small fraction of so-called "high-risk" uses of AI, like credit scoring, sorting software for recruitment, verification of travel documents, robot assisted surgery, the management of critical infrastructure (e.g., electricity), or when an AI assists a judicial authority, to name a few practical examples.

The binding requirements ensure that in such cases, high-quality data sets are used, risks are adequately managed, documentation and logs are kept, and human oversight is provided for – in order to ensure the AI systems are robust, secure, and accurate.

In the end of the day, the purpose of the regulation is two-fold: (a) to ensure that Europeans can trust what AI has to offer and embrace AI-based solutions with confidence they are safe, while (b) to encourage innovation to develop in an ecosystem of trust. As the European Commission's Executive Vice President Vestager put it recently: "Trust is a must, not a 'nice to have.'"

As was the case with the General Data Protection Regulation, the Commission's AI proposal will be subject to legislative scrutiny before it can become law in all EU countries.

Transatlantic and Multilateral Context

Without a doubt, it will also be a topic of some debate here in the United States, as well. And the European Union looks forward to these discussions with our like-minded partners.

This is because, on the global stage, AI has become an area of strategic importance, at the crossroads of geopolitics and security. Having taken this pioneering step, the EU will work to deepen partnerships, coalitions and alliances with third countries and with likeminded partners to promote trustworthy, ethical AI. Exploring a Social Contract for the AI Age – a framework to ensure an AI "Bill of Rights" in the digital age – is fundamental in international relations today.

And in this work, our relationship with the United States is paramount. For Europe and the United States in particular, our shared values make us natural partners in the face of rival systems of digital governance. Together, we must rise to the occasion.

That is why President Ursula von der Leyen called for a Trans-Atlantic Agreement on AI that protects human dignity, individual rights, and democratic principles, to also serve as a "blueprint" for broader global outreach.

My hope is that Europe and the United States will work more closely together, continuously and at all levels – with engineers, policymakers, thought leaders, civil society, scientists, and businesses on both sides of the Atlantic – to guide our technological progress and help us improve, evolve, and become more just, equitable, and free societies. To help us ensure that AI enhances the human condition and experience for all mankind.

And I rely on all AI students and researchers, innovators, policymakers, and business leaders listening in today, to help us turn this aspiration into our new reality.

We have a window of opportunity to act – and we should do so. During my time as EU Ambassador to the United States, I will do all in my power to bring it about.

Once again, please accept my warm gratitude and appreciation for this award.

To view the full event where the speech was given:
https://bostonglobalforum.org/news-and-events/events/artificialintelligence-internationalaccord/

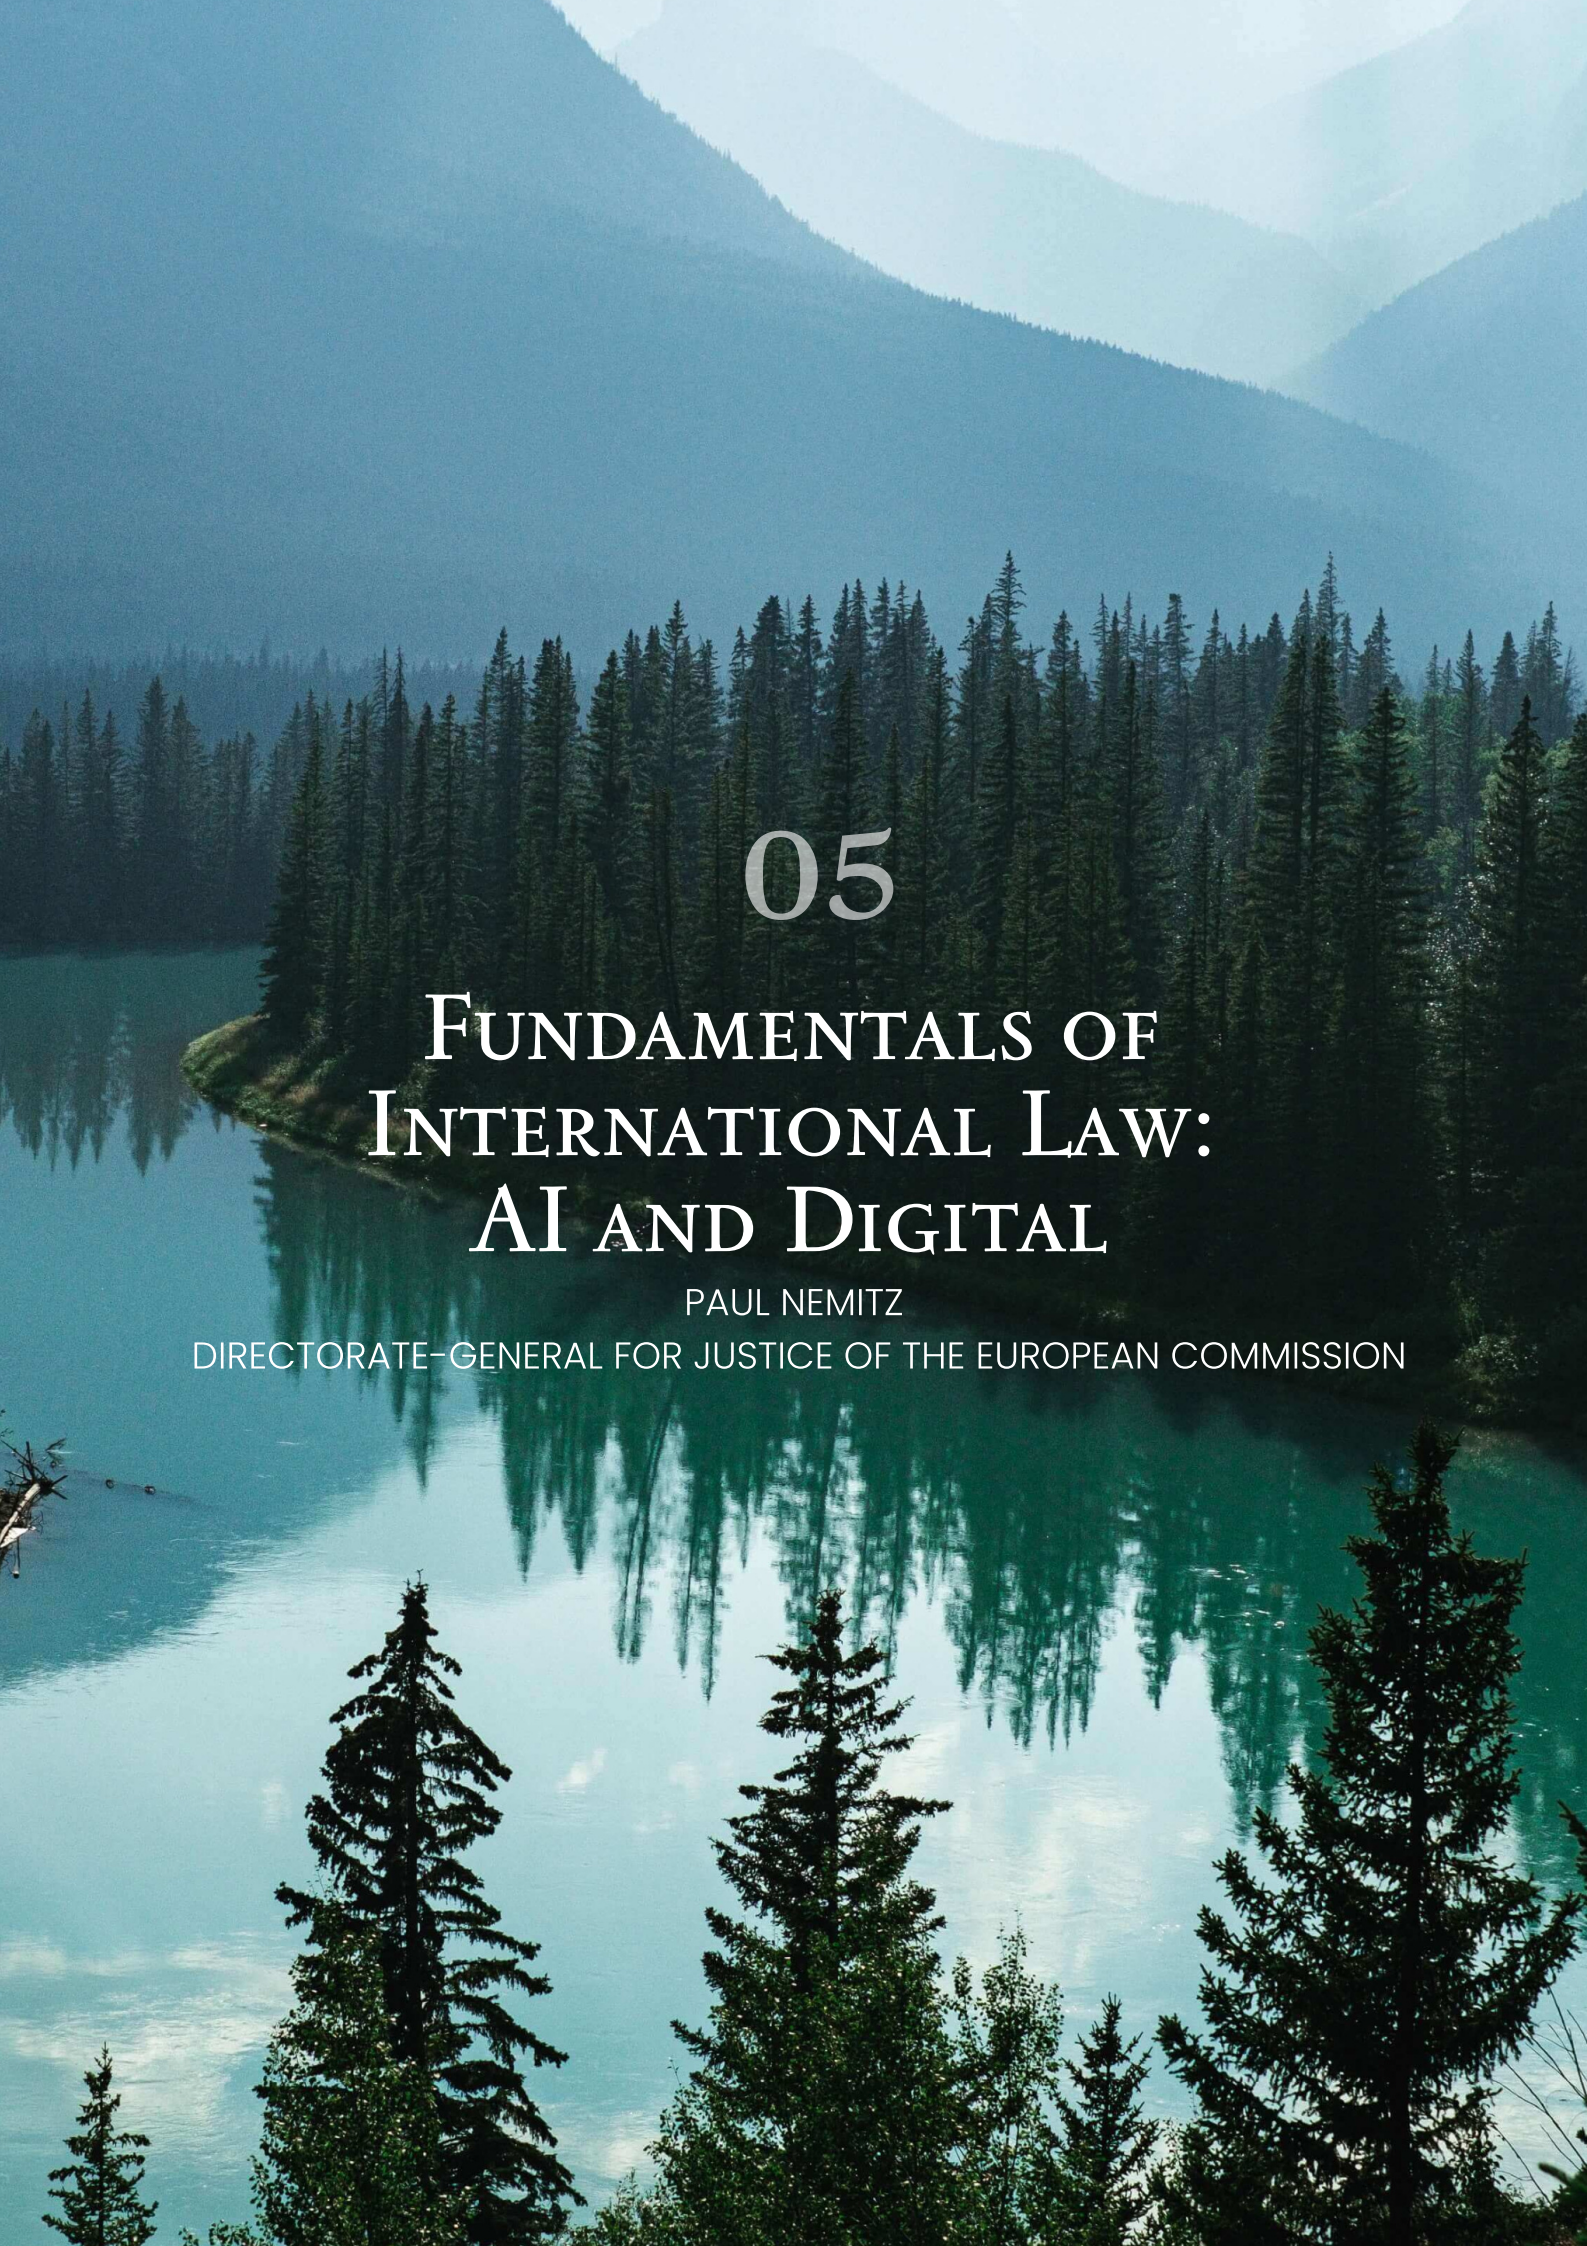

05

Fundamentals of International Law: AI and Digital

PAUL NEMITZ
DIRECTORATE-GENERAL FOR JUSTICE OF THE EUROPEAN COMMISSION

05

Fundamentals of International Law on AI

Paul Nemitz
Directorate-General for JUSTICE of the
European Commission

(Remarks of Paul Nemitz[1] at the Boston Global Forum Conference, April 28, 2021)

Let me give a few words of introduction of possible thinking on a global accord on Artificial Intelligence (AI).

First, because AI is predicted to be pervasive in all areas of society as a useful tool, whether in the sustainable management for and of the environment, education, health, police, justice and security and so on, one may be tempted to rewrite all the rules for society in such an accord. That of course then leads to the necessary tolerance and relativism which we find in the few very broad international agreements which reflect the different views of states on issues such as Human Rights and Democracy. That would be the classic approach to respect all the diversities of countries' views on how things should be going in this form of democracy or that form of democracy.

But we have an alternative to proceeding in this way in relation to AI due to the technical nature of the subject we are dealing with. What do I mean by this? The technical nature of the subject opens up parameters and categories on which I believe agreement is possible, very independent of the question of whether the signatory is a democracy or not and whether a state has this or that security or economic interest or this or that fundamental right in place or not. We should first search for those principles, which will be useful to all in the world in order to be able to agree on binding rules. I will give you some examples. For example, because of the technical nature of AI, we want to be certain in the future that an AI system really does what it is supposed to do. I would say that this is one of those generic principles, which is on a sufficient level of abstraction that everybody in this world could agree to it. In addition, it still would be a useful rule. Another principle would be that AI should never be developed or put on the market for the purpose of undermining the governability of States. So what I would encourage us to do in our ongoing work is to search together for those generic principles which are actually defining what is new in this international accord on AI rather than trying to rewrite all the general principles, ranging from the balance between security and fundamental rights right through to sustainability in environmental affairs, of all policy areas in which AI may be used or of all world accords which may have existed in history. I think we should focus on the new challenge, the technological nature of the challenge, and the new specificities of AI compared to previous challenges of humanity. My prediction is if we do this, we will find principles on which agreement in a binding form is more likely than if we choose to work on general principles on which the world community has only in very rare lucky historic occasions found the possibility to agree on binding rules.

My second remark is relating to the reason why one would want to have such a binding agreement. In the European Union we have moved on from a debate on Ethics for AI to a proposed legal regulation of AI by a parliamentary law. There was really a development of thinking moving from first the thinking that the law is not the right tool and binding rules are not amenable to be used for regulating such a technology. This was the early thinking which led the European Commission to set up a high-level group on AI Ethics. Also, the German government set up a high-level group on data ethics. All in all, we have today around 80 ethics codes around the world on AI. I don't want to speculate on who has put the money on this orientation on ethics rather than the law, for whom it was good and for whom it maybe wasn't, but one thing is clear. In the European Union the thinking has moved on from the conviction that ethics will be at the center as a tool to a recognition that AI needs rules which are binding for all, have democratic legitimacy and therefore also can be enforced against all, using public power. Of course, in terms of substantive principles we want to act ethically. But the Union has moved on to make a start to work on binding rules which are enforceable against everyone. This is significant. Why is this significant? Because we do believe that AI is so important in terms of its ubiquity but also in terms of the power it gives to the state and to private parties that we cannot only work with incentives, non-binding rules and the instruments we know from the OECD, such as ethics, self-regulation and non-binding recommendations. We already have an OECD code on AI principles to which many countries have agreed. We now have the work ongoing in UNESCO. We have work ongoing on some kind of global instrument in the Council of Europe which is the broader European organization to which Russia and Turkey belong. On occasion, the Council of Europe has actually produced binding rules, such as the Budapest convention on Cybercrime to which the United States has also signed up. I would say now we should be ambitious in our work also on AI. In terms of what we aim for, we see that we aim for a binding international agreement of the type of the United Nations Convention on the Law of the Sea (UNCLOS) which contains important elements of a binding nature and which includes also the instruments and institutions which are necessary to make sure that such binding rules are respected by all, such as a Court.

This will be my third remark. Is it right to take the multilateral agreement on the Law of the sea which is a United Nations agreement as a model for our work on AI? The multilateral agreement on the law of the sea is the biggest binding convention of the United Nations. I have in mind the comments which we heard from our first speaker, namely the ambition to have as many countries on board as possible as a precondition for such an accord working. I think that makes the law of the sea convention (UNCLOS) a great model. UNCLOS at the time was extremely innovative in dealing with a space which was largely unexplored, out there beyond the sight from the coast, 70 percent of the space on earth. States agreed on the rules of UNCLOS to ensure peace in this space and to allow an orderly and sustainable use of this space, the sea, by the world community. This is exactly the reason why I believe states should also be able and willing to agree on global binding rules for AI, namely to make it possible to use this technology in an orderly and sustainable manner, and in a way that contributes to peace on earth. The rules of UNCLOS are based on amazingly simple principles. They are binding in nature. They agreed because everybody understood that this new frontier, the oceans being at the beginning of exploitation by deep sea mining and so on, needs common rules to avoid conflict in the future. The United Nations and states went very far in this, even set up a court and a dispute settlement system. Pacem in Maribus, peace on the seas, was the motto, which got the negotiations on UNCLOS going, in the middle of the cold war. As at the time, I think we should have our ambitions high. We should go beyond just repeating non-binding types of principles, which we already find in the OECD code and in many ethics codes. We should focus on the specific nature of the technology which provides us with the opportunity to set out principles on which agreement is possible. We should aim for binding rules because we are dealing with power relations here. Power relations between states and private corporations. And in relation to individuals. Let us not forget that those corporations which are at the forefront of operationalizing AI, in terms of financial capacity, are bigger than many states of this world. So, I can imagine that many states actually have an interest in binding rules for this reason alone. We are dealing with power relations between the states and individuals, not only individuals of their own states. And power relations between corporations which use AI and people. These power relations and the risks involved are one of the key reasons why we are dealing in a binding way with AI in legislation in the European Union. We should recognize that this element of managing power from technology and controlling power from technology is one of the key challenges of the future. It underlies our willingness to come together in this group and look for global principles which help us to give a frame to this new power of technology, of AI, over people, over states, over resources, over security.

To close I recommend to everyone the new report from Bruce Schneier at the Belfer Center at Harvard University that came out this week, on AI as a hacking tool in the future.[2] He describes a world in which all binding principles, all decency, all systems of security which we may have in technical devices will be systematically undermined by AI systems which have been built for this purpose. So I do not believe that we need an agreement which becomes a marketing tool for AI. I do believe we underestimate the risks of AI. And I do believe here we have another principle which we will find agreement on by many. Namely, that AI which has this "purpose to undermine" as described by Bruce Schneier should be ruled out, namely AI which serves as a tool for criminal purposes or to undermine governability of states.

---

1. *The author expresses a person's opinion and not necessarily that of the European Commission*
2. *The coming AI Hackers, April 2021, https://www.belfercenter.org/publication/coming-ai-hackers*

To view the full event at which the remarks were delivered: https://bostonglobalforum.org/news-and-events/events/artificialintelligence-internationalaccord/

06 AIWS 7-Layer model (BGF-G7 Summit Report 2018)

BOSTON GLOBAL FORUM

The model was developed by Boston Global Forum's Michael Dukakis, Nazli Choucri, Allan Cytryn, Alex Jones, Tuan Anh Nguyen, Thomas Patterson, Derek Reveron & David Silbersweig.